



NAVAL POSTGRADUATE SCHOOL

MONTEREY, CALIFORNIA

THESIS

COUNTER PIRACY: A REPEATED GAME WITH ASYMMETRIC INFORMATION

by

Christopher D. Marsh

September 2009

Thesis Advisor:
Second Reader:

Kyle Y. Lin
Timothy H. Chung

Approved for public release; distribution is unlimited

THIS PAGE INTENTIONALLY LEFT BLANK

REPORT DOCUMENTATION PAGE			<i>Form Approved OMB No. 0704-0188</i>	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instruction, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington DC 20503.				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE September 2009	3. REPORT TYPE AND DATES COVERED Master's Thesis	
4. TITLE AND SUBTITLE Counter Piracy: A Repeated Game with Asymmetric Information			5. FUNDING NUMBERS	
6. AUTHOR(S) Christopher D. Marsh				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Postgraduate School Monterey, CA 93943-5000			8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING /MONITORING AGENCY NAME(S) AND ADDRESS(ES) OPNAV N20			10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES The views expressed in this thesis are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government.				
12a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution is unlimited			12b. DISTRIBUTION CODE	
13. ABSTRACT (maximum 200 words) <p>This thesis presents a model of a counter-piracy operation, where a task force has one operational asset (a destroyer) and one reconnaissance asset (an unmanned aerial vehicle) to reduce piracy in a large region. The region is divided into small areas, and each day the pirates operate in one area to hijack commercial vessels to collect ransoms. The information is asymmetric to the two players. The pirates know which area is more profitable, but the task force does not. The task force can use the operational asset to prevent piracy, and the reconnaissance asset to collect information on the profitability of each area. The pirates want to maximize their income over a thirty-day period, while the task force wants to minimize it. The numerical experiments quantify the value of the operational asset and the reconnaissance asset in this counter-piracy operation.</p>				
14. SUBJECT TERMS Piracy, Game Theory, Bayesian update			15. NUMBER OF PAGES 57	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UU	

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std. Z39-18

THIS PAGE INTENTIONALLY LEFT BLANK

Approved for public release; distribution is unlimited

**COUNTER PIRACY: A REPEATED GAME WITH ASYMMETRIC
INFORMATION**

Christopher D. Marsh
Commander, United States Navy
B.S., United States Naval Academy, 1993

Submitted in partial fulfillment of the
requirements for the degree of

MASTER OF SCIENCE IN OPERATIONS RESEARCH

from the

**NAVAL POSTGRADUATE SCHOOL
September 2009**

Author: Christopher D. Marsh

Approved by: Kyle Y. Lin
Thesis Advisor

Timothy H. Chung
Second Reader

Robert F. Dell
Chairman, Department of Operations Research

THIS PAGE INTENTIONALLY LEFT BLANK

ABSTRACT

This thesis presents a model of a counter-piracy operation, where a task force has one operational asset (a destroyer) and one reconnaissance asset (an unmanned aerial vehicle) to reduce piracy in a large region. The region is divided into small areas, and each day the pirates operate in one area to hijack commercial vessels to collect ransoms. The information is asymmetric to the two players. The pirates know which area is more profitable, but the task force does not. The task force can use the operational asset to prevent piracy, and the reconnaissance asset to collect information on the profitability of each area. The pirates want to maximize their income over a thirty-day period, while the task force wants to minimize it. The numerical experiments quantify the value of the operational asset and the reconnaissance asset in this counter-piracy operation.

THIS PAGE INTENTIONALLY LEFT BLANK

TABLE OF CONTENTS

I.	INTRODUCTION.....	1
A.	BACKGROUND ON PIRACY AND COUNTER PIRACY OPERATIONS.....	1
	1. Overview on Piracy.....	1
	2. Operations to Suppress Piracy	2
	3. Challenges and Threats Resulting from Piracy in Somalia	3
	4. Current U.S. Policy and Limitations.....	4
	5. Joint and Navy Doctrine to Implement Policy	4
B.	RESEARCH OBJECTIVES.....	6
C.	RELATED LITERATURE.....	6
	1. Counter-Piracy Models	6
	2. Game Theory and Search Theory	6
D.	SCOPE, LIMITATIONS, AND ASSUMPTIONS	7
II.	THE MODEL	9
A.	SCENARIOS	9
B.	MATHEMATICAL MODEL	12
C.	STRATEGIES	14
	1. Task Force Strategy.....	14
	2. Pirate Strategy.....	18
III.	NUMERICAL DEMONSTRATIONS AND ANALYSES	19
A.	REWARD FOR THE PIRATES	20
B.	LEARNING FOR THE TASK FORCE	24
C.	DISCUSSION	30
IV.	CONCLUSION AND RECOMMENDATIONS.....	33
A.	MATHEMATICAL LIMITATIONS AND ASSUMPTIONS.....	33
B.	FUTURE STUDIES.....	33
	LIST OF REFERENCES.....	37
	INITIAL DISTRIBUTION LIST	39

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF FIGURES

Figure 1.	Depiction of the region	9
Figure 2.	IMB depiction of pirate activity.....	11
Figure 3.	Expected daily reward for the Pirates through the season for Scenario 1. Maximum standard error in cases 2 and 3 after 1000 simulations is less than \$1,000 throughout the 30 day season.....	21
Figure 4.	Expected daily reward for the Pirates through the season for Scenario 2. Maximum standard error in cases 2 and 3 after 1000 simulations is less than \$2,000 throughout the 30 day season.....	22
Figure 5.	Task Force perception the probability the state of nature is the true state in Scenario 1. Maximum standard error in cases 2 and 3 after 1000 simulations is less than .008 over the 30 day season.	27
Figure 6.	Task Force perception the probability the state of nature is the true state of nature in Scenario 2. Maximum standard error in cases 2 and 3 after 1000 simulations is less than .008 over the 30 day season.	28
Figure 7.	Days required for the Task Force to determine the probability the true state of nature is greater than 90% for a given scenario and case.....	30

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF TABLES

Table 1.	Sample pirate reward matrix.....	19
Table 2.	The value added by Task Force's assets from Figure 3 and 4.....	23
Table 3.	Days required to obtain knowledge of the actual state of nature.....	25

THIS PAGE INTENTIONALLY LEFT BLANK

EXECUTIVE SUMMARY

Optimizing intelligence collection with limited resources is a common problem for operational commanders. The dilemma facing operational commanders is how to reconcile the conflict between maximum intelligence collection and maximum operational effects. This thesis presents a counter-piracy model with two scenarios as an example of the conflict between the effects of operational and reconnaissance assets.

The scenarios represent small scale operations with a Task Force that has one Destroyer and one Unmanned Aerial Vehicle (UAV) to prevent a group of Pirates from hijacking commercial vessels. The region where the Pirates operate is broken into three areas. The Destroyer can prevent the Pirates from operating in one area each day, while the UAV collects information about one area each day. The reward the Pirates receive for hijacking a vessel in an area varies according to merchant traffic density and environmental conditions. The Pirates are familiar with the local environment and know the reward distribution for operating in each area, while the Task Force does not. The Task Force employs the Destroyer to deter hijackings and can learn the reward distributions to maximize the effects of the Destroyer.

The model compares the reward the Pirates receive over a thirty-day period and the time required for the Task Force to learn the true state of nature in four cases. The cases are (1) when the Task Force has one Destroyer that cannot collect intelligence, (2) when the Destroyer can collect intelligence, (3) when a UAV that can collect intelligence is added, and (4) when the Task Force knows the true state of nature. The scenarios are further divided into three simulations with different variances in the Pirate's reward.

The numerical experiments show the Pirate's reward decreasing significantly as the amount of reconnaissance is increased through each case. The results also show increased effects of reconnaissance when variance in Pirate reward increases. The model and the numerical experiments provide insight into tasking methodologies for reconnaissance and operational platforms.

THIS PAGE INTENTIONALLY LEFT BLANK

ACKNOWLEDGMENTS

I would like to thank my advisor, Kyle Lin, and second reader, Timothy Chung, for their guidance, patience, and support throughout this process. Kyle Lin spent countless hours instructing and guiding me to create the scenario and model through several iterations.

I would also like to thank my sponsor the Maritime Operating Center and Maritime Headquarters directorate at Fleet Forces Command and the staff at the Johns Hopkins Applied Physics Lab for their time and support in researching thesis topics. In particular Jack Keane, Steve Carr, and John Fletcher were very helpful in researching possible topics.

Finally, I would also like to thank my family and friends who have been understanding and patient with me throughout my time at NPS. Without the humorous comments and sympathetic ears, this thesis would not have been possible.

THIS PAGE INTENTIONALLY LEFT BLANK

I. INTRODUCTION

A. BACKGROUND ON PIRACY AND COUNTER PIRACY OPERATIONS

1. Overview on Piracy

Piracy occurs in nearly every maritime realm and has been a threat to legal commerce for thousands of years. Bertrand Russell cites the early reasons for piracy in the Mediterranean in his History of Western Philosophy.

Weapons, until about 1000 B.C., were made of bronze, and nations which did not have the necessary metals on their own territory were obliged to obtain them by trade or piracy. Piracy was a temporary expedient, and where social and political conditions were fairly stable, commerce was found to be more profitable. (Russell, 1972)

Outside the Mediterranean, several regions are famous for piracy. Particularly the Caribbean in the 17th and 18th centuries, the Barbary Pirates of the same era, the straits of Malacca from 14th century to present, and most recently in the news is piracy off the coast of Somalia. The common thread is that piracy was more profitable than legitimate commerce for a variety of reasons including a lack of natural resources, abundance of valuable trade routes, easy access to weapons, and a lack of governments to provide maritime security. Modern piracy takes many forms, such as robbery of vessels at sea or at anchor, the hijacking of vessels at sea, and kidnap for ransom attacks (Raymond, 2009).

Combating piracy requires several aspects to decrease the allure of piracy. Peter Leeson, a noted economist at George Mason and author of The Invisible Hook: The Hidden Economics of Pirates, was quoted in the New York Times blog stating, “We have to recognize that pirates are rational economic actors and that piracy is an occupational choice. If we think of them as irrational, or as pursuing other ends, we’re liable to come up with solutions to the pirate problem that are ineffective.” (Hagen, 2009)

2. Operations to Suppress Piracy

Naval counter-piracy operations take many forms including escort operations through high risk sea lanes, naval presence operations, and direct assault against pirate land bases. The success or failure of these operations depend on the environment, type of pirates, nature of commercial targets, and resources available to the counter-piracy forces.

The most famous counter-piracy example in U.S. history is the attack on the Barbary Pirates in the early 19th century. A land force of Americans and Arabs on the outskirts of Tripoli, produced a peace treaty with sponsors of regional pirates, signed on June 5, 1805. Although this treaty included tribute of \$60,000, it was attributed to a change in philosophy of European governments on their policy of tribute. The era of terror and crime on the high seas in North Africa was over (Turner, 2003).

Efforts in the Straits of Malacca, a long time hot spot for piracy, is another example of counter-piracy operations conducted by regional navies to establish legal procedures. In 1992, the International Maritime Bureau created the Piracy Reporting Center in Kuala Lumpur. The reporting center brought attention to the regional problem and combined with the threat of terrorism to require action from regional partners. International pressures from the U.S. and Japan particularly helped encourage Malaysia and Indonesia to work with the Singapore Navy in coordinated patrols of the region. Increased cooperation in the region includes the agreement for the Information Sharing Center in Singapore for fourteen nations. Combined with increased regional stability in the Aceh Province of Indonesia piracy was reduced significantly after 2004 (Raymond, 2009). These efforts highlight the importance in combined efforts to make piracy physically difficult while removing the underlying cause by facilitating more profitable enterprises in the region.

The Gulf of Aden represents one aspect of Somali piracy with unique issues. The Gulf of Aden is a highly trafficked region with several unstable states around it, particularly Yemen and Somalia. The high density of merchant traffic in the constrained space make easy targets for pirates. A recent proliferation of piracy in the region in 2008 brought forth significant international cooperation in the form of naval task forces from

several countries including Russia, India, and China, as well as members of the Coalition Naval Forces in the U.S. Central Command area of operations. The U.S. stood up TASK FORCE 151 to coordinate the patrols in the region. The concentration of naval assets in the constrained area brought several successes in the form of foiled attacks. During the summer and fall of 2008, two dozen attacks were repelled by U.S. FIFTH Fleet warships (Hassan & Gutterman, 2008). Piracy still occurs in the region, and as of the spring of 2009, 250 mariners and dozens of ships were being held for ransom. The International Maritime Organization sponsored a meeting in January of 2009 to coordinate efforts of regional nations to develop a coherent approach. The strategy is reflected in the Djibouti Code of Conduct. The Djibouti Code focuses on building the diplomatic, legal, and military capabilities of the regional nations including Somalia to be able to counter all aspects of piracy.

Despite the successful examples of counter-piracy operations, future and ongoing crises will be constrained by resources. Military and diplomatic leaders must compromise on the amount of support they can offer. The demand will continually exist for a combined set of metrics legally, militarily, economically, and diplomatically to prevent piracy.

3. Challenges and Threats Resulting from Piracy in Somalia

The problem with piracy in the Gulf of Aden is significantly different than the problem in the Somali Basin. The vast expanse of water, combined with the large number of fishing villages on the Somali east coast, prevent effective saturation by naval forces. Large transit distances prevent escort operations. The co-location of pirate's bases of operations with fishing villages inhibits military strikes on pirate bases. The instability of the Somali government and the fractured tribal structure of the fishing villages further complicate the problem and prevent diplomatic or economic solutions. U.S. Agency for International Development, through their famine early warning network, notes Somalia's increased reliance on foreign foods arriving in Somali ports and the associated decrease in regional stability. The threat of piracy further increases commodity prices, decreases income in commercial trade, and delays shipments

throughout the region (U.S. Agency for International Development, 2009). The result is a cycle that increases the incentive for Somali's to turn to piracy and decreases legitimate commercial incentives. The threat of famine increases risk of piracy during relief operations. When international government organizations and nongovernment organizations attempt to send relief supplies, pirates can hijack supplies and increase their profits and local prestige.

4. Current U.S. Policy and Limitations

While the U.S. is pursuing a combined policy that combines State and Defense Department resources in accordance with the Djibouti Code of Conduct, the challenges of Somalia are daunting. The interagency response, referred to as a Maritime Operational Threat Response (MOTR) plan works with the International Maritime Bureau and regional partners to encourage conditions that discourage piracy in the region (Kraska & Wilson, 2009). However, the training of the Somali Coast Guard is focused on the more lucrative area of the Gulf of Aden instead of the larger region of the Somali basin. Recent initiative, such as the Djibouti Code of Conduct, will improve local conditions and encourage lawful behavior, but change will take time. The lack of infrastructure, complex tribal organizations, and vast length of the Somali east coast guarantee progress will move slowly. This leads to the question of how much we can accomplish with a small military force operating in a large region where pirates operate.

5. Joint and Navy Doctrine to Implement Policy

The implementation of the national policy requires guidance for operational employment. Joint publications provide the guidance required to develop operational measures of effectiveness and intelligence requirements. Joint publications also include guidance for measures of performance. Tactical guidance requires documents specific to individual platforms. Operational guidance is contained within Service and Combatant Commander's guidance. These are typically derived from the overriding publications from the Joint Chiefs of Staff.

Doctrine related to the operational employment of reconnaissance in support of a task force is contained in Joint Publication 2.01.3 Joint Tactics Techniques and Procedures for Intelligence Preparation of the Battlespace.

The primary purpose of reconnaissance is to gain information to facilitate the JIPB [Joint Intelligence Preparation of the Battlespace] support to the operational level is concerned with analyzing the operational area, facilitating the flow of friendly forces in a timely manner, sustaining those forces, and then integrating tactical capabilities at the decisive time and place. (Joint Chiefs of Staff, 2000)

This document is the primary source for understanding the flow of information during operational planning and provides guidance on the development of intelligence requirements.

Joint Operational Planning Joint Publication 5.0 is the primary document for operational planners to assist in understanding the operational environment and developing operational effects. Combining effects with the understanding of the operational environment is critical to successful planning. This paper attempts to identify a model to fulfill the operational planning guidance contained.

Commanders continuously assess the operational environment and the progress of operations, and compare them to their initial vision and intent. The assessment process begins during mission analysis when the commander and staff consider what to measure and how to measure it to determine progress toward accomplishing a task, creating an effect, or achieving an objective. Commanders adjust operations based on their assessment to ensure objectives are met and the military end state is achieved. (Joint Chiefs of Staff, 2006)

While the guidance for planning intelligence requirements and operational effects are contained within the publications, the formulation is up to the field commanders. Specific metrics to connect the intelligence requirements and operational effects are developed intuitively and often lack specific measures of performance or effectiveness that can be readily analyzed.

B. RESEARCH OBJECTIVES

Much of the research on piracy focuses on the tactical approach of interdiction and capture of pirates and their vessels. This thesis is intended to address operational issues that face commanders when allocated few resources to patrol large regions. Problems of how to allocate the resources and equate operational objectives with intelligence collection are part of all military operations. Without a common metric to determine operational effects and intelligence collection, it is impossible to adequately allocate the scarce resources. This thesis explores one possible approach to identifying a common metric for the effects of operations and intelligence.

C. RELATED LITERATURE

1. Counter-Piracy Models

The most comprehensive counter-piracy model is the model produced by the Naval Postgraduate School Systems Engineering Analysis Department for the Straits of Malacca in their 2005 report *Maritime Domain Protection in the Straits of Malacca*. This model incorporates a five module simulation including sensors, command and control, force models, land inspections and sea inspections. This model focused on reducing attacks while minimizing operational costs and impacts to regional commerce. The model produced exhaustive reports on potential threats to regional shipping, cost benefit analysis of operational assets, and analysis of regional commerce. (Systems Engineering Analysis Cohort Seven, 2005)

Other counter-piracy models focus on the ability to identify and interdict pirates through maritime interdiction operations. These models use queuing theory to maximize the number of ships that can be searched in a given region. These models are often used when trying to clear a smaller region from known threats as in the case of studies to support Task Force 151 escort operations.

2. Game Theory and Search Theory

Because pirates and counter-piracy forces have opposing goals, it is natural to use game theory to analyze the piracy problem. One problem in the application of game

theory to military operations is the ability to accurately quantify a payout matrix in the face of uncertainty. The basic problem of identifying units to associate with the payout matrix usually results in probabilities as in anti submarine warfare and ratio of forces in Melvin Dresher's "Tactical Air War Game," (Dresher, 1961).

Payout matrices still have the problem of uncertainty. Several solutions to problems with uncertainty have been produced over the years, but two stand out. First the work of John Harsanyi in developing games with incomplete information identified the information available to each player as a type in a Bayesian Game (Myerson, 2004). This work also demonstrated examples of how to exploit an opponent's erroneous beliefs and an explanation on complications resulting from the normal form of the game. The work of Robert Aumann and Michael Maschler tackled the problem of repeated games with a lack of information and developed a solution methodology that influenced this model (Maschler, 1995).

D. SCOPE, LIMITATIONS, AND ASSUMPTIONS

The scope of this paper is intended to address the operational allocation problems faced by a small task force. For this reason, several assumptions are required to focus the research on the desired problem. The primary assumptions are in the capabilities of the platforms. The platforms are given the ability to accurately observe several variables and determine a singular accurate value. This does not account for several problems in reconnaissance that include false detections. This model also assumes the Pirates are interested only in monetary reward. Sources indicate this is true to a degree, but the complexity of criminal organizations and the regional tribal structure are not accounted for in the model. The model assumes a single entity in control of piracy within the region. This model is limited to scenarios where the interests of two parties are directly opposed resulting in a two-person zero-sum game.

THIS PAGE INTENTIONALLY LEFT BLANK

II. THE MODEL

This chapter discusses the modeling effort. In Section A, we introduce the scenarios and motivations of our models. In Section B, we define the mathematical models. In Section C, we discuss the strategies that we want to study for both the Pirates and the Task Force.

A. SCENARIOS

The scenarios model simple counter-piracy operations off the coast of Somalia, where a small number of ships are assigned to patrol a large region against Pirates targeting commercial vessels for hijacking and ransom. The region is divided into several small areas (an example with three areas is shown in Figure 1). The Pirates operate in one area each day. A Task Force, equipped with one Destroyer and one Unmanned Aerial Vehicle (UAV), is assigned to deter the Pirate's operation and to protect the region. The Task Force cannot see the Pirates, who blend into the local fishing fleet, but can prevent a Pirate's attack in an area with the presence of a Destroyer. At the dawn of each day, the Pirates select an area to operate during the day, while the Task Force decides where to allocate the Destroyer and the UAV. These daily operations are repeated each day for a season, while the Pirates attempt to earn as much as possible and the Task Force attempts to minimize the Pirates earnings.




Figure 1. Depiction of the region

The Pirates' expected reward is obtained from recent studies on the economics of pirates (McIntyre, 2009) and data from the International Maritime Bureau (ICC International Maritime Bureau, 2009). The Pirates' expected reward is estimated between \$400,000 and \$800,000 per day during peak seasons. This range is based on the assumptions that the Pirates capture between six and eight ships per month and collect a ransom of between one and three million dollars per ship. Operating costs, due to the cost of boats, weapons, and the care and feeding of the Pirates, is assumed to be negligible compared to the estimated profit. The gangs of pirates are estimated to contain about 1,000 people. The Pirate crews collect significantly more than the average Somali yearly income, which is about \$600 per year. Variations in ransom from the capture of a vessel includes uncertainties caused by the merchant vessels unwillingness to reveal actual ransom amounts, costs of negotiators, and delivery costs the pirates assume.

The Pirates focus operations in the area that gives them the highest rewards based on the number of commercial ships operating and the ease of capturing them in that area. Despite increased cooperation between commercial vessels and counter-piracy forces, merchant vessels still travel through warning areas, as evidenced by the number of ships attacked off of Somalia this year (ICC International Maritime Bureau, 2009). Figure 2 shows the attack locations off the Somali coast between January and June of 2009.



Total attacks Gulf of Aden / southern Red Sea / east coast Oman and Arabian Sea – 103
Total attacks east coast Somalia and Indian Ocean - 45

Figure 2. IMB depiction of pirate activity

The reward the Pirates earn from operating in an area varies according to merchant vessel routing, sea states, and weather conditions, and is modeled by a normal distribution. The mean of the reward distribution is between \$400,000 and \$800,000 while the standard deviation is between \$100,000 and \$200,000. The novelty of our model is that the Pirates, which consist of local gangs, have more information about these variations than the Task Force. Consequently, the Pirates know precisely the distribution of rewards by operating in each area. The Task Force, on the other hand, knows some areas are more profitable than the others, but the Task Force does not know precisely which area is the most (or least) profitable. Specifically, we consider two scenarios as follows: In the first scenario, commercial vessels avoid the central regions by cutting corners transiting to Kenya or the Gulf of Aden as they pass through the outer areas. The low density of commercial vessels in the central area results in a consistently low expected reward. The outer areas have more vessel traffic and contain higher expected rewards. Depending on the local conditions one outer area is easier for the Pirates to

operate, therefore more profitable than the other outer area. The Pirates know which area is the most profitable one, but the Task Force does not.

In the second scenario, the bulk of the merchant traffic travels through the central area with variations on the outer areas. This situation is common during humanitarian relief efforts, when there is a high density traffic route to one of the neighboring ports. In this situation, the expected reward in the central area contains the highest reward. The variances occur on the fringe of the traffic route. The local conditions make one of the outer areas more difficult to operate, hence less profitable than the other outer area. While the Pirates have complete information about each area's value, the Task Force knows the center area is most profitable but does not know which outer area is least profitable.

In both scenarios the Task Force can learn about the state of nature by operating in the outer areas, but not by operating in the center area. The contrast between the two scenarios represents differences in operational allocation problems. The first scenario represents a problem where the operational and intelligence collection requirements are aligned with each other. In this scenario, the Task Force can gain most information by operating in the areas with the largest reward to the Pirates. The second scenario represents a conflict between operations and intelligence. In this case, preventing piracy in the most profitable areas does not provide any information about the actual state of nature.

B. MATHEMATICAL MODEL

Suppose the whole region is divided into I small areas. Each day, the Pirates select one area to operate in, while the Task Force selects one area to send its Destroyer. The planning horizon consists of T days in a season, during which the Pirates want to maximize the expected total reward, while the Task Force wants to minimize the total reward. There are K possible states of nature. For state of nature k , the Pirates know the mean $\mu_{i,k}$ and the standard deviation $\sigma_{i,k}$ of the reward, if the Pirates operate in area i . The Pirates learn the actual state of nature k^* at the beginning of each season, but the Task Force does not and has to initially assume that each state of nature is equally likely. The Task Force attempts to minimize the reward of the Pirates by choosing a mixed

strategy defined by the probabilities of operating in each area. This Task Force game is produced by a weighted average of reward matrices in all states of nature. The result is referred to as the average game.

To assess the values of the Task Force's assets, we consider four cases as follows:

1. The Task Force has one Destroyer, which does not have any surveillance capability. The Task Force assigns the Destroyer to operate in one area at the beginning of the day. If the Pirates and the Destroyer occupy the same area, the Pirates will observe the Destroyer and not hijack any vessels that day and receive no reward. If the Pirates and Destroyer choose different areas, the Pirates will hijack a vessel and receive the reward for the chosen area. Because the Destroyer does not have any surveillance capability, the Task Force does not learn about the true state of nature and continues to play the average game introduced on the first day of the season.
2. The Task Force has one Destroyer, which has surveillance capability. The Destroyer conducts surveillance on the environment while protecting commercial vessels from the Pirates attack in one area. The surveillance collected is transformed into a single number that represents the reward if the Pirates operate in that area without the presence of the Destroyer. The Task Force daily allocation is made according to the mixed strategy corresponding to the average game, as in the previous case. The only difference is that the Task Force can update the probability on the state of nature each day.
3. The Task Force has one Destroyer and one UAV, both of which have surveillance capability. The UAV has no ability to deter the Pirates, but can collect information on the area. The allocation of the UAV is made after the Task Force determines the location of the Destroyer. The UAV is sent to an area that the Destroyer does not occupy and provides information about the state of nature. If the Destroyer goes to one outer area, the UAV goes to the other. If the Destroyer goes to the central area, then the UAV is randomly assigned to one outer area with probability 0.5.

4. The Task Force learns the state of nature before the season begins. The mixed strategy employed by the Task Force is the optimal mixed strategy to minimize the Pirates' reward in the matrix representing the true state of nature. This case is used as a benchmark to assess the value of the Task Force's surveillance capability.

The Pirates' operations are hidden within the local fishing fleets and are not visible to the Task Force. After a hijack, the Task Force knows of the incident but does not learn the reward or area of the hijacking. The only information the Task Force can gain about the state of nature is the information about the region they operate in on a specific day. The Pirates know about the Task Force's lack of information and apply a pure strategy that maximizes their reward against the Task Force mixed strategy.

The Task Force attempt to minimize the reward includes efforts to learn the actual state of nature. This creates a common problem between deploying assets to perform an operational mission vice a reconnaissance mission. The reconnaissance mission can learn about the state of nature and improve the mixed strategy the Task Force uses, but if the Task Force only has one Destroyer, then the reconnaissance mission reduces the immediate operational effect. The Task Force can overcome this through the allocation of a separate reconnaissance platform such as a UAV to operate independently of the Destroyer and learn the true state of nature. In each scenario, the two states of nature are symmetric so that the value of each state played as a game with a mixed strategy will have equal values.

C. STRATEGIES

1. Task Force Strategy

The Task Force strategy is considered a myopic strategy because it uses information available on day t to minimize the Pirates' reward on day $t+1$ without taking into account how the learning on day $t+1$ might affect the future reward. With the myopic strategy, the Task Force first computes the average game between two possible states using the updated state probabilities. The Task Force then computes the optimal mixed strategy in this average game. This produces the myopic value, which is also equal

to the value of the game, if no further information is collected. The process of collecting information about the environment determines the Task Force's perception about the true state of nature. The Task Force's perception is represented by the probability $p_k(t)$ a given state k is the true state of nature at a given day t . The Task Force's perception is updated after collecting information about an area.

The update of $p_k(t)$ is conducted through observation each day operations are conducted in an area. Consider the case when the Task Force has one Destroyer. The initial belief of the Task Force is that each state of nature is equally likely. After the Destroyer occupies an area j for one day, it observes the local conditions for that day and observes the reward $r_j(t)$. The observed reward varies day-to-day according to the distribution representing the expected reward in the area. The Task Force can then compute an updated probability that the state of nature is each of the K possible states. We assume the reward follows a normal distribution, with the following density function where μ is the mean and σ the standard deviation.

$$f(x, \mu, \sigma) = \frac{e^{-(x-\mu)^2/(2\sigma^2)}}{\sqrt{2\pi}\sigma}$$

The process of collecting information is modeled using a Bayesian update. In Case 1, no information is gained from the Destroyer and the values of $p_k(t)$ remain constant for all t . In Case 2 (the Destroyer collects information on the area) and Case 3 (the Destroyer and UAV collect information), the Task Force learns about an area in the form of the observed reward $r_j(t)$ in area j for time t to update $p_k(t)$. If the Task Force has complete information, as in Case 4, no update is required as the value of $p_k(t)$ is equal to one when k is equal to k^* and equal to zero otherwise. The information gained about one area is then used to update the Task Force beliefs about the state of nature using Bayes Law using the probabilities $p_k(t-1)$ from the prior day.

$$p_k(t) = \frac{[p_k(t-1)]f(x, \mu_{i,k}, \sigma_{i,k})}{\sum_{k=1}^K [p_k(t-1)]f(x, \mu_{i,k}, \sigma_{i,k})}$$

The Task Force then computes the next day's average game, and uses it to determine a new mixed strategy.

The Task Force strategy is a vector of probabilities $y_i(t)$ over the possible operating areas for a given day t . The strategy is chosen to minimize the reward, or game value $v_k(t)$ in state k at stage t . This game value uses the weighted average reward $\bar{\mu}_i(t)$ of the average game computed by the weighted average of the corresponding state's matrix $\mu_{i,k}$, with the weight equal to the Task Force perception of a state of nature $p_k(t)$. The computation to determine the reward the Task Force expects begins with the following linear program to solve the value of the average game.

FORMULATION (G1) :

$$\min_y \bar{v}(t) \tag{G1.1}$$

$$\text{Such that } (1 - y_i(t))\bar{\mu}_i(t) - \bar{v}(t) \leq 0 \quad \forall i \tag{G1.2}$$

$$\sum_{i=1}^I y_i(t) = 1 \tag{G1.3}$$

$$\text{Where: } \bar{\mu}_i(t) = \sum_{k=1}^K \mu_{i,k} p_k(t) \tag{G1.4}$$

Since, $\bar{\mu}_i(t) = 0$ for some i , then $y_i(t) = 0$. In other words, if the Pirates cannot collect any reward in area i , then the Task Force does not need to send the Destroyer to area i . If $\bar{\mu}_i(t) > 0$ then in the optimal solution the constraint in Equation (G1.2) is tight. In other words, the equality holds in Equation (G1.2), which yields

$$y_i(t) = 1 - \frac{\bar{v}(t)}{\bar{\mu}_i(t)} \tag{G1.6}$$

This provides an analytical method to determine optimal employment, if the value of the game is known. While the value of the game is unknown at this point, we know that $y_i(t)$ is a probability, and the sum of $y_i(t)$ over the set i is equal to one. This implies that

$$\sum_{i=1}^I (1 - \frac{\bar{v}(t)}{\bar{\mu}_i(t)}) = 1 \quad (\text{G1.7})$$

Summing over the set of possible strategies I this can be simplified into

$$I - \bar{v}(t) \sum_{i=1}^I \frac{1}{\bar{\mu}_i(t)} = 1 \quad (\text{G1.8})$$

This leads to an analytical result for the value of the average game in terms of the number of areas in the region and the rewards for each area.

$$\bar{v}(t) = \frac{I - 1}{\sum_{i=1}^I \frac{1}{\bar{\mu}_i(t)}} \quad (\text{G1.9})$$

While the Task Force's strategy for the next day is computed using equation (G1.6) and (G1.9), an update is conducted to compute $p_k(t)$ and identify the actual state of nature.

Originally, we described the learning process with only the Destroyer. Next, we consider the case when the Task Force also has a UAV and can use it to gain information about the state of nature. The UAV, once assigned to an area on a given day, observes the reward value in that area for the day, but does not deter the Pirates' attack. The difference between the two states of nature is based on the difference between the two outer areas in each scenario. The updated information on the reward value for outside regions is helpful in learning the true state of nature. Therefore, if the Destroyer goes to an outside area, it is optimal to assign the UAV to the other outside area. If the Destroyer goes to the central area, then we assign the UAV to each outside area with probability 0.5.

2. Pirate Strategy

The Pirates are familiar with the region. In order to assess the value of each Task Force asset, we consider a worst-case scenario by assuming that the Pirates are able to observe the action taken by the Task Force on a daily basis. Therefore, the Pirates know what the Task Force learned about the area and apply the same Bayes formula to predict Task Force's mixed strategy on the next day. Consequently, the Pirates can apply the best pure strategy against Task Force's mixed strategy each day. The familiarity of the region allows the Pirates to fully capitalize on the lack of information on the side of the Task Force. The resulting reward computed in G2 is greater than what the Task Force expects, and can be viewed as a worst-case scenario from Task Force's standpoint.

The expected reward the Pirates can receive $\hat{r}(t)$ is computed using a pure strategy against the Task Force mixed strategy $y_j(t)$. The reward is computed for a given day t by:

FORMULATION (G2) :

$$\hat{r}(t) = \max_{i \in I} \mu_{i,k^*} (1 - y_i(t)) \quad (\text{G2})$$

The Pirates update the Task Force perception of the states of nature and computed the reward value and pure strategy every day prior to sending out their boats. The pure strategy is the optimal strategy the Pirates can employ knowing the true state of nature, while the Task Force uses the myopic strategy based on the average game. The Pirates only change their behavior based on the perceptions of the Task Force. The Pirates do not change their behavior based on the allocation of the UAV.

III. NUMERICAL DEMONSTRATIONS AND ANALYSES

We implemented the model in a simulation using Microsoft Excel with Visual Basic for Applications. In each scenario, we consider four cases that represent different Task Force capabilities. The scenarios are different in the estimated rewards the Pirates receive by operating in each area. The mean value of the rewards can be \$400K, \$600K, and \$800K, but the primary difference is the location of the known and unknown values. For each scenario, we vary the standard deviation of reward among \$100K, \$150K, or \$200K.

Table 1. Sample pirate reward matrix.

State k= 1		Task Force Strategies		
		Area 1	Area 2	Area 3
Pirate Strategies	Area 1	0,0	$\mu_{1,1}$, $\sigma_{1,1}$	$\mu_{1,1}$, $\sigma_{1,1}$
	Area 2	$\mu_{2,1}$, $\sigma_{2,1}$	0,0	$\mu_{2,1}$, $\sigma_{2,1}$
	Area 3	$\mu_{3,1}$, $\sigma_{3,1}$	$\mu_{3,1}$, $\sigma_{3,1}$	0,0

From Table 1 the reward matrices for each scenario represent one of the two states of nature. In Scenario 1, the values for $\mu_{1,1}$, $\mu_{2,1}$, and $\mu_{3,1}$ are \$600K, \$400K, and \$800K, respectively and the values for $\mu_{1,2}$, $\mu_{2,2}$, and $\mu_{3,2}$ are \$800K, \$400K, and \$600K, respectively. Scenario 2 sets the values of $\mu_{1,1}$, $\mu_{2,1}$, and $\mu_{3,1}$ at \$600K, \$800K, and \$400K and $\mu_{1,2}$, $\mu_{2,2}$, and $\mu_{3,2}$ at \$400K, \$800K, and \$600K.

The simulations ran 1000 times for each level of standard deviation. Without loss of generality, we set the true state to be state one, because of the symmetry between two states. We consider two measures of effectiveness: (1) the cumulative reward the Pirates receive over the season and (2) the number of days required for the Task Force to learn the probability state one is the true state of nature is greater than 90%.

A. REWARD FOR THE PIRATES

We compute the cumulative reward by summing over the Pirates' expected daily reward for the duration of the season. The Pirates maximize this value by choosing the best pure strategy against the Task Force's myopic strategy on each day.

To derive the values of different assets of the Task Force, we consider the four cases discussed in Chapter II. In Case 1, the Task Force sends the Destroyer into the region, without collecting any information, but prevents the Pirates from operating freely. Case 1 represents the operational effect of the Destroyer. Case 2, which allows the Destroyer to collect information, represents the combined operational and reconnaissance effect of the Destroyer. Case 3 represents the effect of the additional reconnaissance provided by a UAV. Case 4, when the Task Force has full information about the true state of nature represents the operational effect with full information. The following graphs depict the daily expected reward of the Pirates. The top and bottom lines represent Case 1 and Case 4. These lines form the upper and lower bounds of the Pirates' daily reward. The areas under the curves represent the cumulative reward values. The areas between the curves represent the benefit of additional capabilities to the Task Force. The areas between Case 1 and Case 2 represent the Pirates' reduced reward due to the surveillance capability by the Destroyer. The areas between Case 2 and Case 3 represent the reduced reward due to additional information gained by the UAV. The areas between Case 3 and Case 4 represent potential reduction in Pirates' reward if the Task Force has full information about the three areas.

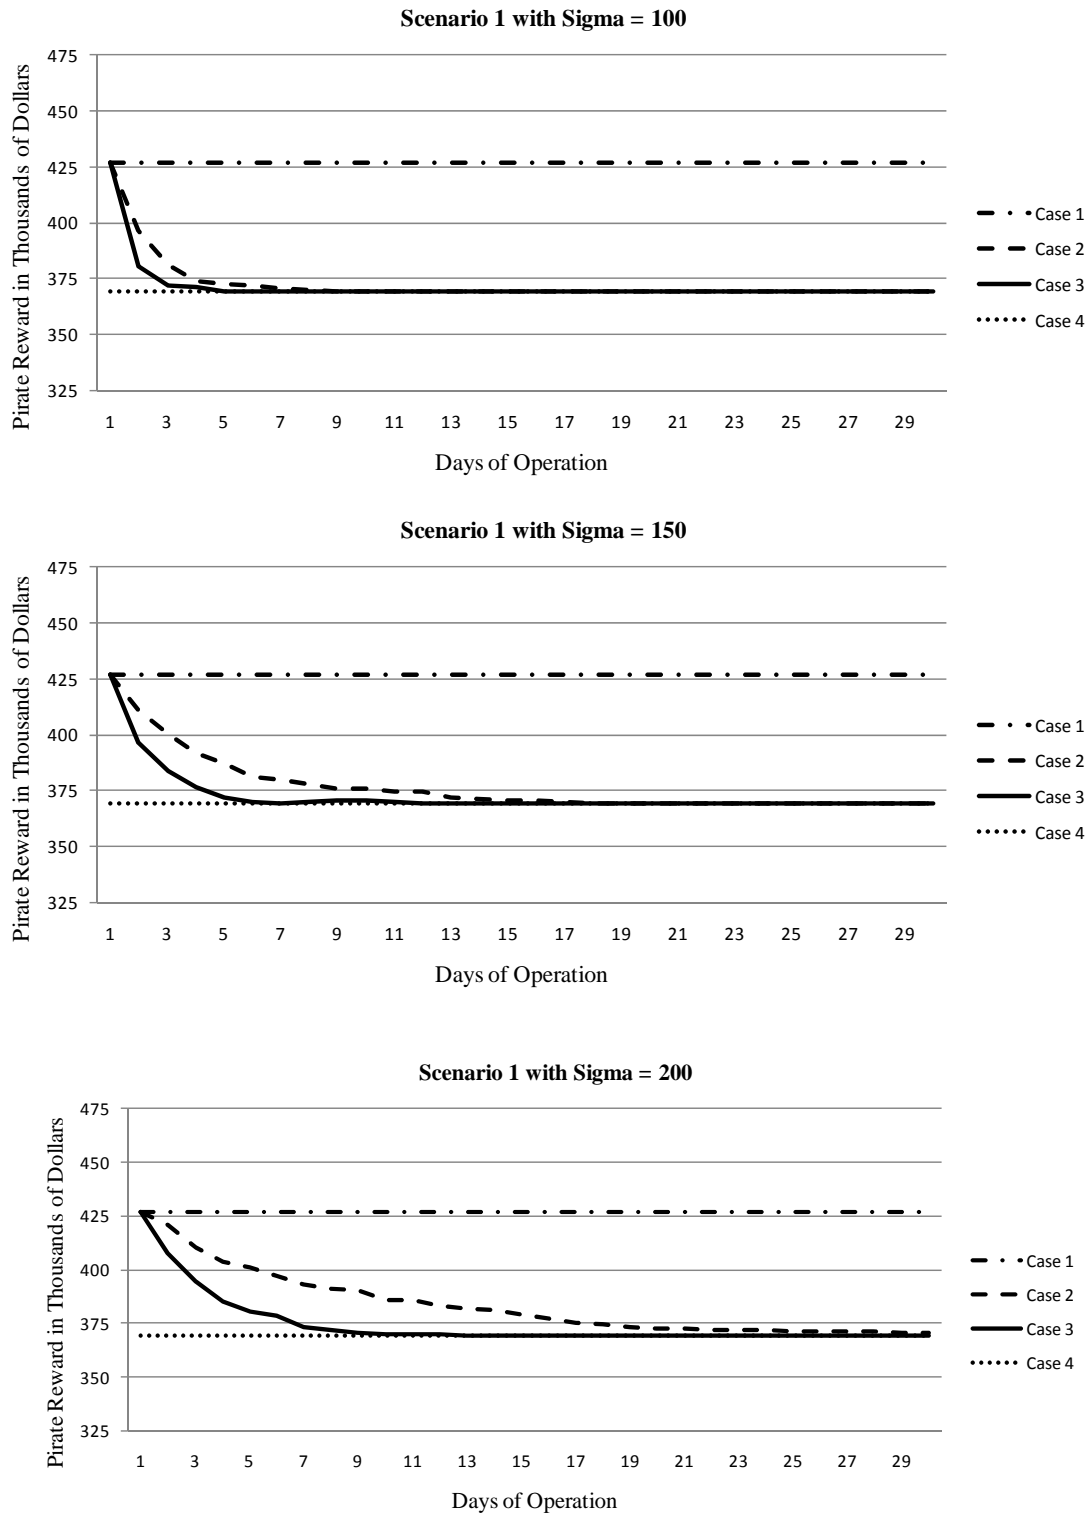


Figure 3. Expected daily reward for the Pirates through the season for Scenario 1. Maximum standard error in cases 2 and 3 after 1000 simulations is less than \$1,000 throughout the 30 day season.

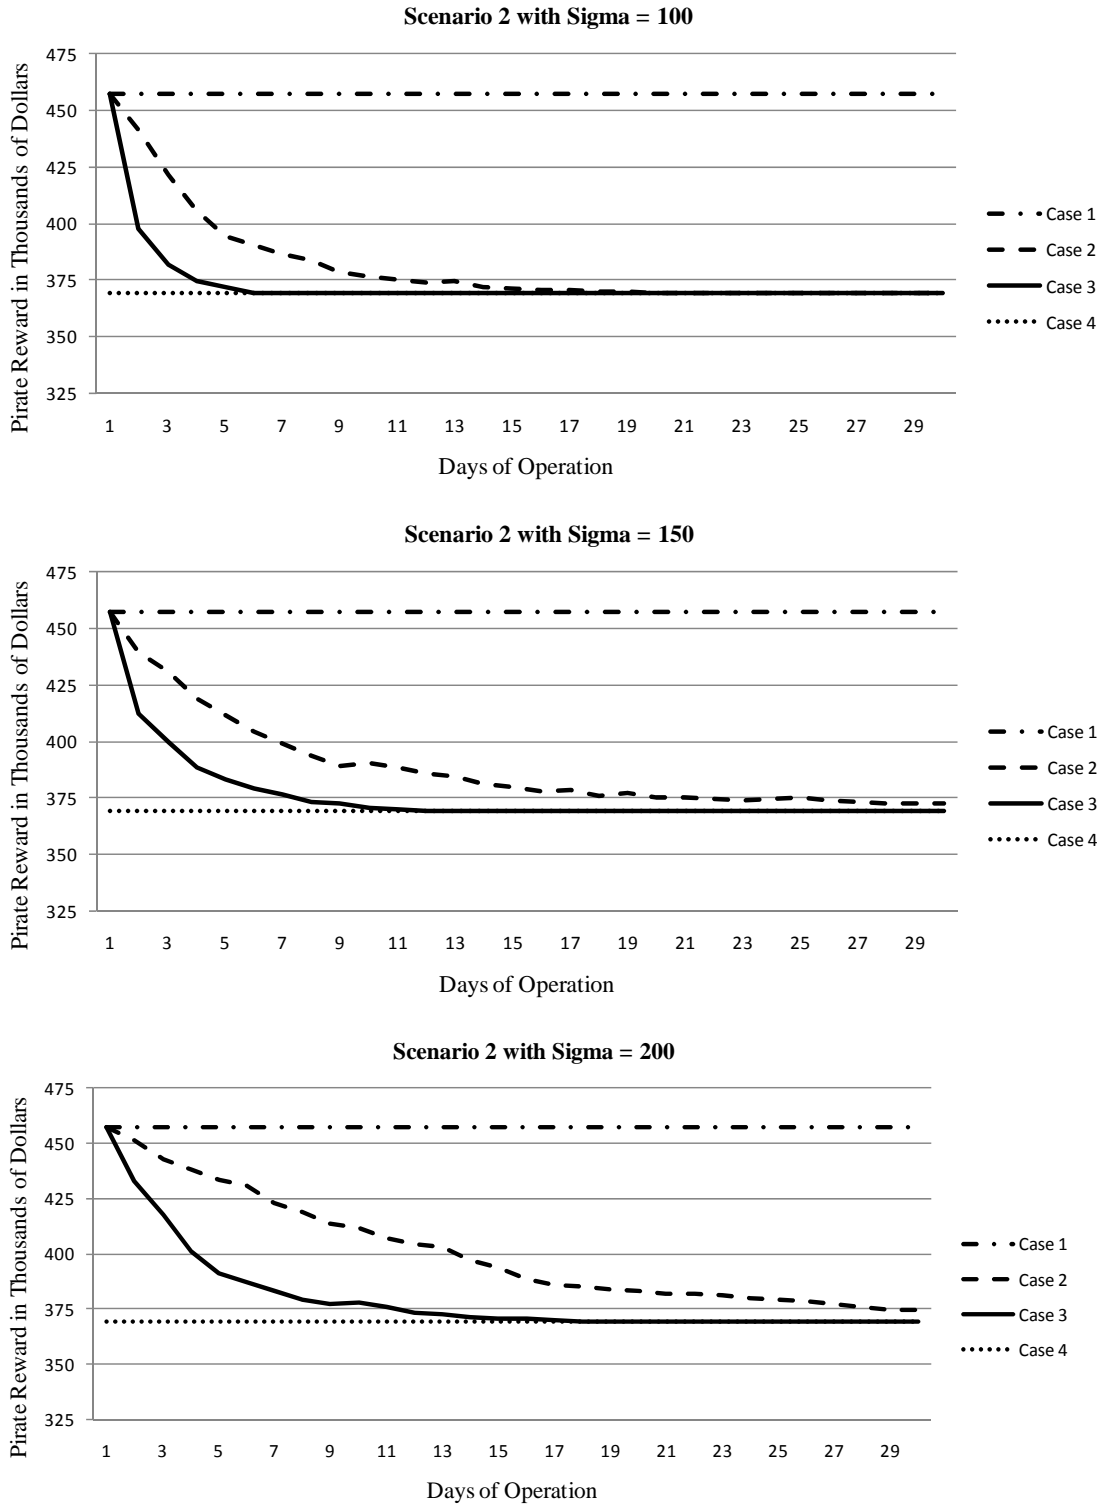


Figure 4. Expected daily reward for the Pirates through the season for Scenario 2. Maximum standard error in cases 2 and 3 after 1000 simulations is less than \$2,000 throughout the 30 day season.

The values from the graphs in Figure 3 and Figure 4 are summarized in the following table.

Table 2. The value added by Task Force's assets from Figure 3 and 4.

		Reduction in Pirate Cumulative Reward (in thousands of dollars)			
Scenario	Sigma	Case 1 DDG w/o ISR	Case 2 DDG w/ISR no UAV	Case 3 DDG w/ ISR and UAV	Case 4 Full Information
Scenario 1	100	11200	1614	36	73
	150	11200	1490	120	114
	200	11200	1282	272	168
Scenario 2	100	10286	2269	203	165
	150	10286	2038	379	221
	200	10286	1678	627	332

Table 2 shows the greatest decrease in Pirates' expected reward is due to the deterrence capability of the Destroyer. Decreasing marginal utility is evident for the information gained by the Destroyer as σ decreases where the Destroyer can learn about the true state of nature more quickly. While the marginal utility of the information from the UAV increases with an increase in σ .

Note that each scenario has a different cumulative reward for Case 1 despite using the same range of reward values. The resultant values may even counter the Task Force's operational intuition. In Scenario 1, where the Task Force does not know which area is most valuable to the Pirates, is the Task Force must spread their one asset across the possible areas to gain the maximum effect. In Scenario 2, the Task Force does know the most valuable area and affects a significant result with a strategy that concentrates on the most valuable area. It could be easy to believe that the uninformed mixed strategy against Scenario 2 would be more effective than the uninformed mixed strategy against Scenario 1. The graphs show the difference in expected reward is actually larger in Scenario 2. This is due to the Pirates taking advantage of the lack of information held by the Task Force. The result is demonstrated by computing the value of the average game for each scenario.

The values for the reward associated with each area are described Chapter II Section A and derived from data from the International Maritime Bureau. In Scenario 1,

if the Destroyer does not have any surveillance capability, then the Task Force plays the average game with the following payout matrix.

$$.5 \begin{bmatrix} 0 & 600 & 600 \\ 400 & 0 & 400 \\ 800 & 800 & 0 \end{bmatrix} + .5 \begin{bmatrix} 0 & 800 & 800 \\ 400 & 0 & 400 \\ 600 & 600 & 0 \end{bmatrix} = \begin{bmatrix} 0 & 700 & 700 \\ 400 & 0 & 400 \\ 700 & 700 & 0 \end{bmatrix}$$

The value of this average game is 373.

The payout matrix for the average game in Scenario 2 follows and has a value of 380.

$$.5 \begin{bmatrix} 0 & 600 & 600 \\ 800 & 0 & 800 \\ 400 & 400 & 0 \end{bmatrix} + .5 \begin{bmatrix} 0 & 400 & 400 \\ 800 & 0 & 800 \\ 600 & 600 & 0 \end{bmatrix} = \begin{bmatrix} 0 & 500 & 500 \\ 800 & 0 & 800 \\ 500 & 500 & 0 \end{bmatrix}$$

It is easy to see that in Scenario 2, there is an increase in the value of the game over Scenario 1 despite using the same numbers. The difference is further exacerbated when the Pirates are allowed to capitalize on the lack of information with a pure strategy as is evidenced by the line representing Case 1 from graphs in Figures 3 and 4.

Table 2 shows the largest decrease in cumulative reward is due to deterrence provided by the presence of the Destroyer. The decrease is constant with respect to σ but does vary with the scenario. As discussed earlier the two scenarios should have different cumulative reward values based on the location of the highest reward area. This is evident by comparing the difference between the curves representing no information gained and complete information on the side of the Task Force. Figures 3 and 4 show the Pirates' daily reward decays toward the same value in each scenario. Differences in the cumulative values in Table 2 are caused by a slower learning process in scenario one. The learning process will be discussed in the following section.

B. LEARNING FOR THE TASK FORCE

The learning process of the Task Force is defined as the ability of the Task Force to learn the true state of nature. The measure of effectiveness is the number of days required for the Task Force to learn the probability state one is the true state of nature is

greater than 90%. In each scenario, the Task Force was able to reliably achieve this goal within the thirty-day season. Still, the longer the true state of nature was ambiguous the more reward the Pirates accumulated.

The previous section detailed a slower learning process derived from observations on cumulative Pirate reward. One reason for the slower learning process is as follows. Given the average game for Scenario 2 is represented by the following matrix:

$$.5 \begin{bmatrix} 0 & 600 & 600 \\ 800 & 0 & 800 \\ 400 & 400 & 0 \end{bmatrix} + .5 \begin{bmatrix} 0 & 400 & 400 \\ 800 & 0 & 800 \\ 600 & 600 & 0 \end{bmatrix} = \begin{bmatrix} 0 & 500 & 500 \\ 800 & 0 & 800 \\ 500 & 500 & 0 \end{bmatrix}$$

It is evident that the Task Force would want to initially use a strategy that focuses on area two to minimize the reward of the Pirates. This slows the learning process because the Destroyer spends most of the time in the area that does not help identify the state of nature. This logic captures the dilemma of allocating assets to maximize operational effects vice maximizing intelligence collection.

The graphs in Figures 5 and 6 demonstrate the Task Force learning process in the scenarios as a measure of probability the state of nature is state one versus the number of days of operations. Table 3 represents a summation of the data in Figures 5 and 6.

Table 3. Days required to obtain knowledge of the actual state of nature.
Days Required Before Probability State of
Nature is State 1 is Greater Than 90%

Scenario	Sigma	Case 2 No UAV	Case 3 W/ UAV
Scenario 1	100	1.94	1.30
	150	4.57	2.23
	200	8.45	3.93
Scenario 2	100	4.35	1.86
	150	8.58	3.40
	200	11.92	5.09

The variation of results by σ is expected due to the difficulty in gathering information with increased uncertainty. The differences between the scenarios contain additional differences. The differences may be more visible through the graph depicting the rate at which information is collected in each case in the following figures.




Figure 5. Task Force perception the probability the state of nature is the true state in Scenario 1. Maximum standard error in cases 2 and 3 after 1000 simulations is less than .008 over the 30 day season.




Figure 6. Task Force perception the probability the state of nature is the true state of nature in Scenario 2. Maximum standard error in cases 2 and 3 after 1000 simulations is less than .008 over the 30 day season.

The different rates of learning are evident in the graphs by noting the difference in area under the curve for scenario two with and without the UAV (Case 2 and Case 3 respectively). The difference in learning rate is further exaggerated when there is greater uncertainty in the individual area represented by σ .

The primary factors affecting the learning process were the assets allocated, the standard deviation of the area, and the scenario. While the number of assets and standard deviation are expected to impact the Task Force ability to learn, the effect of the scenario requires further analysis. The Task Force myopic strategy focuses on the area they believe is most valuable. In Scenario 1, this is not as significant because the area with the highest reward according to the average game is an area that contains information about the true state of nature. The result is that efforts to maximize the operational effect will also maximize the rate the Task Force learns the true state of nature.

Scenario 2 highlights a dilemma in tasking operational and intelligence platforms. The most effective Task Force mixed strategy for the average game focuses on Area 2 because it provides the least reward to the Pirates. Unfortunately, the Task Force cannot learn about the true state of nature by operating in Area 2, since it is the same in both states of nature. The differences are most pronounced in Case 2 when there is no UAV to focus on intelligence collection. Figure 7 demonstrates how this effect is made more prominent as σ increases.




Figure 7. Days required for the Task Force to determine the probability the true state of nature is greater than 90% for a given scenario and case.

The differences can also be seen in the curves of Case 2 in Figures 5 and 6. The numbers in Table 3 confirm this as well.

C. DISCUSSION

Two observations may help improve allocation of operational assets. The first involves an understanding of the tactical employment of an operational asset through a game theoretic perspective. The second accounts for the value of information in a scenario.

The first observation about the employment of the Destroyer in Scenario 2 concerns the allocation to a non-informative area. The most obvious method to avoid this problem is to focus initial allocation to areas that provide information about the state of nature. The addition of the UAV solves this problem because it always operates in an informative area. Simulation is required to determine the optimal number of days before

reverting to a purely myopic strategy. For example, in Scenario 2, without a UAV available, a commander may choose to focus on intelligence collection for 12 days prior to maximizing operational effects.

The second observation is that the use of collecting information reaches a point of diminishing returns. The utility of the UAV decreases over time and is apparent with the converging values of Case 2 and Case 3 in Figures 4 through 7. If the Task Force objective is to change behavior patterns by decreasing the Pirate's cumulative reward, continuous reconnaissance may not be required. However, this model does not account for search factors that further degrade the ability to collect information. This will be discussed further in recommendations.

THIS PAGE INTENTIONALLY LEFT BLANK

IV. CONCLUSION AND RECOMMENDATIONS

A. MATHEMATICAL LIMITATIONS AND ASSUMPTIONS

The limits of the model are divided into two aspects. First, the assumptions of the game theoretic construct using the assumptions based on behavioral aspects will be examined. Second, the limitations of the model will be analyzed for computational efficiency.

The two-person, zero-sum game theoretic construct requires the model be limited to two players with diametrically opposed rewards. One player is assumed to have complete information and the other player with some predisposed belief. This assumption allows the model to function as a two-person zero-sum game with a lack of information on one side. One additional assumption is complete information is available to one side, which is not always the case. The uncertainty in the information available to the Pirates was modeled by using a normal distribution to represent the reward value for the Pirates. The result permitted the Pirates to act as if they had perfect information.

The ability to expand the model in terms of areas within the region, strategies of the players, and possible states of nature can be accomplished with some cost in the amount of computation required. The formulations are called once per turn of the simulation. One additional consideration is the time required to run individual simulations and the duration of the season.

B. FUTURE STUDIES

The model has the potential to be expanded for future use by using more complex game theoretic constructs, incorporating actual sensor data, or incorporating more detailed models of the reward functions. The advantage of each is to increase the accuracy and detail of the model. Some expansions of the model also have the potential to model different aspects of conflict including information warfare, coalition building, and intelligence analysis.

Nonzero-sum games, vice the current zero-sum game, have the potential to model more complex scenarios where the interests of the players are not diametrically opposed. The nonzero-sum game would allow a more diverse scenario and application into operations that are not specifically designed to counter a specific enemy action.

Since most counter-piracy operations are coalition efforts, there is a benefit to incorporate n-person games to understand the dynamics and potential rewards to be gained through a coalition. More than two players in a game create significant complications, but can yield information relating to the effectiveness added by individual coalition members. This would benefit coalition building efforts by helping to determine command structures and incentives offered by coalition leaders. Guillermo Owen, in his book Game Theory discusses several examples of coalition games that could incorporate a lack of information into the reward structure (Owen, 1995).

The model developed in this thesis could also be applied in the context of information warfare. Specifically, instead of the Task Force using a learning process to gain information about the state of nature, the Pirates could send disinformation to deny the Task Force access to the actual state of information. This could also include a lack of information on each side, where both sides participate in a learning or disinformation process. The process would require additional simulations, and the Pirates would have to adopt the myopic strategy as well.

Incorporating actual sensor data from platforms would offer the opportunity to study the effects of false indications, imperfect probability of detection, and actual sensor coverage area. One example where this could be useful is to address a common operational dilemma of tasking reconnaissance assets. Reconnaissance assets are often assigned in two ways, direct support or associated support. Direct support assigns the reconnaissance asset to work directly for the operational asset. Associated support assigns the reconnaissance asset to work separately from the operational asset. This model specifically addressed associated support. Some benefits of the direct support are greater area of operational effect, increased accuracy of collection due to sensor fusion, and improved communications between the assets. The advantage of associated support, as in this model, is that operational effects can be maximized without the constraint of

intelligence collection requirements and vice versa. A possible mechanism to address this would be two models each using a different allocation method. The direct support model allows the operational asset to cover a larger space decreasing the number of possible strategies. This model is similar to Case 2. The associated support model allows the assets to cover multiple areas, but the areas are smaller resulting in a greater number of possible strategies.

The assumption that the information collected in a given area is readily translated into a specific reward value from a distribution is very different from the reality of intelligence collection. Intelligence is typically tasked to the reconnaissance asset through a list of requirements that the reconnaissance asset can observe, such as number of ships in an area. The observables that form the essential elements of information are difficult to translate into a specific value. Regression analysis may be a mechanism to translate several variables, such as merchant traffic density, sea state, and weather into a specific reward value. This would provide the opportunity to study the effectiveness of different capabilities against specific elements of information.

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF REFERENCES

- Dresher, M. (1961). Mathematics of Games of Strategy. Santa Monica, California: Rand Corporation.
- Hagen, R. (2009, April 20). "Freakonomics." Retrieved August 24, 2009, from New York Times: <http://freakonomics.blogs.nytimes.com/2009/04/20/pirate-economics-101-a-qa-with-invisible-hook-author-peter-leeson/>
- Hassan, M. O., & Gutterman, S. (2008, November 12). "Russian, British Ships Repel Somali Pirate Attack." Retrieved July 15, 2009, from USA Today: http://www.usatoday.com/news/world/2008-11-12-3512568313_x.htm
- ICC International Maritime Bureau. (2009). Piracy and Armed Robbery Against Ships Report for the Period 1 January - 30 June 2009. London: ICC International Maritime Bureau.
- Joint Chiefs of Staff. (2006). Joint Operational Planning. Washington, DC: Defense Technical Information Center.
- Joint Chiefs of Staff. (2000). Joint Tactics Techniques and Procedures for Intelligence Preparation of the Battlespace. Washington DC: Defense Technical Information Center.
- Kraska, J., & Wilson, B. (2009). "Combating Pirates of the Gulf of Aden: The Djibouti Code and the Somali Coast Guard." *Ocean and Coastal Management* , 964–969.
- Maschler, R. A. (1995). Repeated Games with Incomplete Information. Cambridge, Ma: The MIT Press.
- McIntyre, D. (2009, April 14). "Somali Pirates are Getting Rich: A Look at the Profit Margins," Retrieved July 14, 2009, from 24/7 Wall Street: <http://247wallst.com/2009/04/14/somali-pirates-are-getting-rich-a-look-at-the-profit-margins/>
- Myerson, R. (2004). "Games with Incomplete Information Played by 'Bayesian Players'". *Management Science* , 1804–1817.
- Owen, G. (1995). *Game Theory*. San Diego: Academic Press.
- Raymond, C. Z. (2009). Piracy and Armed Robbery in the Malacca Strait. *Naval War College Review* , 31–42.
- Russell, B. (1972). The History of Western Philosophy. New York: Simon and Schuster.

Systems Engineering Analysis Cohort Seven. (2005). Maritime Domain Protection in the Straits of Malacca. Monterey: Naval Postgraduate School.

Turner, R. F. (2003). "State Responsibility and the War on Terror: The Legacy of Thomas Jefferson and the Barbary Pirates." *Chicago Journal of International Law*, 121–140.

U.S. Agency for International Development. (2009, April). "The impact of piracy on livelihoods and food security in Somalia." Retrieved July 15, 2009, from Famine Early Warning Systems Network Home Page:
<http://www.fews.net/docs/Publications/1000872.pdf>

INITIAL DISTRIBUTION LIST

1. Defense Technical Information Center
Ft. Belvoir, Virginia
2. Dudley Knox Library
Naval Postgraduate School
Monterey, California
3. Johns Hopkins Applied Physics Lab
Laurel, Maryland
4. Director of Naval Intelligence
Arlington, Virginia